

On the propriety of restricted Boltzmann machines

Andee Kaplan
Iowa State University
ajkaplan@iastate.edu and
Daniel Nordman
Iowa State University
dnordman@iastate.edu and
Stephen Vardeman
Iowa State University
vardeman@iastate.edu

Abstract

A restricted Boltzmann machine (RBM) is an undirected graphical model constructed for discrete or continuous random variables, with two layers, one hidden and one visible, and no conditional dependency within a layer. In recent years, RBMs have risen to prominence due to their connection to deep learning. By treating a hidden layer of one RBM as the visible layer in a second RBM, a deep architecture can be created. RBMs are thought to thereby have the ability to encode very complex and rich structures in data, making them attractive for supervised learning. However, the generative behavior of RBMs is largely unexplored. In this paper, we discuss the relationship between RBM parameter specification in the binary case and the tendency to undesirable model properties such as degeneracy, instability and uninterpretability. We also describe the difficulties that arise in likelihood-based and Bayes fitting of such (highly flexible) models, especially as Gibbs sampling (quasi-Bayes) methods are often advocated for the RBM model structure.

Keywords: Degeneracy, Instability, Classification, Deep Learning, Graphical Models

1 Introduction

The data mining and machine learning communities have recently shown great interest in deep learning, specifically in stacked restricted Boltzmann machines (RBMs) (see R. Salakhutdinov and Hinton 2009; R. Salakhutdinov and Hinton 2012; Srivastava, Salakhutdinov, and Hinton 2013; Le Roux and Bengio 2008 for examples). A RBM is a probabilistic undirected graphical model (for discrete or continuous random variables) with two layers, one hidden and one visible, with no conditional dependency within a layer (Smolensky 1986). These models have reportedly been used with success in classification of images (Larochelle and Bengio 2008; Srivastava and Salakhutdinov 2012). However, the model properties are largely unexplored in the literature and the commonly cited fitting methodology remains heuristic-based and abstruse (Hinton, Osindero, and Teh 2006). In this paper, we provide steps toward a thorough understanding of the model class and its properties from the perspective of statistical theory, and then explore the possibility of a rigorous fitting methodology. We find the RBM model class to be deficient in two fundamental ways.

First, the models are often unsatisfactory as conceptualizations of how data are generated. Recalling Fisher (1922), the aim of a statistical model is to represent data in a compact way. Neyman and Box further state that a model should “provide an explanation of the mechanism underlying the observed phenomena” (Lehmann 1990; G. E. P. Box 1967). At issue, RBMs often produce data lacking realistic variability and may thereby fail to provide a satisfactory conceptualization of a data generation process. In addition to such degeneracy, we find that RBMs can easily exhibit types of instability. In practice, this may be seen when a single pixel change in an image results in a wildly different classification in an image classification problem. Such model impropriety issues have recently been documented in RBMs (Li 2014), as well as other deep architectures (Szegedy et al. 2013; Nguyen, Yosinski, and Clune 2014). We investigate these phenomena for RBMs in Section 3 through simulations of small, manageable examples.

Secondly, the fitting of these models is problematic. As the size of these models grows, both maximum likelihood and Bayesian methods of fitting quickly become intractable. The literature often suggests Markov chain Monte Carlo (MCMC) tools for approximate

maximization of likelihoods to fit these models (e.g., Gibbs sampling to exploit conditional structure in hidden and visible variables), but little is said about achieving convergence (Hinton 2010; Hinton, Osindero, and Teh 2006). Related to this, these MCMC algorithms require updating potentially many latent variables (hiddens) which can critically influence convergence in MCMC-based likelihood methods.

In Section 4.1, we compare three Bayesian techniques involving MCMC approximations that are computationally more accessible than direct maximum likelihood, which also aim to avoid parts of a RBM parameter space that yield unattractive models. As might be expected, with greater computational complexity comes an increase in fitting accuracy, but at the cost of practical feasibility.

For a RBM model with enough hidden variables, it has been shown that any distribution for the visibles can be approximated arbitrarily well (Le Roux and Bengio 2008; Montufar and Ay 2011; and Montúfar, Rauh, and Ay 2011). However, the empirical distribution of a training set of vectors of visibles is the best fitting model for observed cell data. As a consequence, we find that any fully principled fitting method based on the likelihood for a RBM with enough hidden variables will simply seek to reproduce the (discrete) empirical distribution of a training set. Thus, there can be no “smoothed distribution” achieved in fitting a RBM model of sufficient size with a rigorous likelihood-based method. We are therefore led to be skeptical that any model built using these structures (like a deep Boltzmann machine) can achieve useful prediction or inference in a principled way without intentionally limiting the flexibility of the fitted model.

This paper is structured as follows. Section 2 formally defines the RBM including the joint distribution of hidden and visible variables and explains the model’s connection to deep learning. Additionally, measures of model impropriety and methods of quantifying/detecting it are defined. Section 3 details our explorations into the model behavior and propriety (or lack thereof) of the RBM class. We discuss three Bayesian fitting techniques intended to avoid model impropriety in Section 4 and conclude with a discussion in Section 5. On-line supplementary materials provide proofs for results on RBM parameterizations and data codings described in Section 3.1.2.

2 Background

2.1 Restricted Boltzmann machines

A restricted Boltzmann machine (RBM) is an undirected graphical model specified for discrete or continuous random variables, binary variables being most commonly considered. In this paper, we consider the binary case for concreteness. A RBM architecture has two layers, hidden (\mathcal{H}) and visible (\mathcal{V}), with no dependency connections within a layer. An example of this structure is in Figure 1 with the hidden nodes indicated by gray circles and the visible nodes indicated by white circles.




Figure 1: An example restricted Boltzmann machine (RBM), consisting of two layers, one hidden (\mathcal{H}) and one visible (\mathcal{V}), with no connections within a layer. Hidden nodes are indicated by gray filled circles and the visible nodes indicated by unfilled circles.

A common use for RBMs is to create features for use in classification. For example, binary images can be classified through a process that treats the pixel values as the visible variables v_i in a RBM model (Hinton, Osindero, and Teh 2006).

2.1.1 Joint distribution

Let $\mathbf{x} = (h_1, \dots, h_{n_H}, v_1, \dots, v_{n_V})$ represent the states of the visible and hidden nodes in a RBM for some integers $n_V, n_H \geq 1$. Each single binary random variable, visible or hidden,

will take its values in a common coding set \mathcal{C} , where we allow one of two possibilities for the coding set, $\mathcal{C} = \{0, 1\}$ or $\mathcal{C} = \{-1, 1\}$, with “1” always indicating the “high” value of the variable. While $\mathcal{C} = \{0, 1\}$ may be a natural starting point, we argue later that the coding $\mathcal{C} = \{-1, 1\}$ induces more interpretable model properties for the RBM (cf. Section 3). A standard parametric form for probabilities corresponding to a potential vector of states, $\mathbf{X} = (H_1, \dots, H_{n_H}, V_1, \dots, V_{n_V})$, for the nodes is

$$f_{\boldsymbol{\theta}}(\mathbf{x}) \equiv P_{\boldsymbol{\theta}}(\mathbf{X} = \mathbf{x}) = \frac{\exp \left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j \right)}{\gamma(\boldsymbol{\theta})}, \quad \mathbf{x} \in \mathcal{C}^{n_H+n_V} \quad (1)$$

where $\boldsymbol{\theta} \equiv (\theta_{11}, \dots, \theta_{1n_H}, \dots, \theta_{n_V 1}, \dots, \theta_{n_V n_H}, \theta_{v_1}, \dots, \theta_{v_{n_V}}, \theta_{h_1}, \dots, \theta_{h_{n_H}}) \in \mathbb{R}^{n_V+n_H+n_V*n_H}$ denotes the vector of model parameters and the denominator

$$\gamma(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \mathcal{C}^{n_H+n_V}} \exp \left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j \right)$$

is the normalizing function that ensures the probabilities (1) sum to one. For $\mathbf{x} = (h_1, \dots, h_{n_H}, v_1, \dots, v_{n_V}) \in \mathcal{C}^{n_V+n_H}$ and

$$\mathbf{t}(\mathbf{x}) = (h_1, \dots, h_{n_H}, v_1, \dots, v_{n_V}, v_1 h_1, \dots, v_{n_V} h_{n_H}) \in \mathcal{C}^{n_H+n_V+n_H*n_V}, \quad (2)$$

let $\mathcal{T} = \{\mathbf{t}(\mathbf{x}) : \mathbf{x} \in \mathcal{C}^{n_V+n_H}\} \subset \mathbb{R}^{n_V+n_H+n_V*n_H}$ be the set of possible values for the vector of variables needed to compute probabilities (1) in the model, and write $Q_{\boldsymbol{\theta}}(\mathbf{x}) = \sum_{i=1}^{n_H} \sum_{j=1}^{n_V} \theta_{ij} h_i v_j + \sum_{i=1}^{n_H} \theta_{h_i} h_i + \sum_{j=1}^{n_V} \theta_{v_j} v_j$ for the “neg-potential” function. The RBM model is parameterized by $\boldsymbol{\theta}$ containing two types of parameters, main effects and interaction effects. The main effects parameters ($\{\theta_{v_i}, \theta_{h_j}\}_{i=1, \dots, n_V, j=1, \dots, n_H}$) weight the values of the visible v_i and hidden h_j nodes in probabilities (1) and the interaction effect parameters (θ_{ij}) weight the values of the connections $v_i h_j$, or dependencies, between hidden and visible layers.

Due to the potential size of the model, the normalizing constant $\gamma(\boldsymbol{\theta})$ can be practically impossible to calculate, making simple estimation of the model parameter vector problematic. A kind of Gibbs sampling can be tried (due to the conditional architecture of the RBM,

i.e. visibles given hidden or vice versa). Specifically, the conditional independence of nodes in each layer (given those nodes in the other layer) allows for stepwise simulation of both hidden layers and model parameters (e.g., see the contrastive divergence of Hinton (2002) or Bayes methods in Section 4).

2.1.2 Connection to Deep Learning

RBMs have risen to prominence in recent years due to their connection to deep learning (see Hinton, Osindero, and Teh 2006; R. Salakhutdinov and Hinton 2012; Srivastava, Salakhutdinov, and Hinton 2013 for examples). By stacking multiple layers of RBMs in a deep architecture, proponents of the models claim to produce the ability to learn “internal representations that become increasingly complex, which is considered to be a promising way of solving object and speech recognition problems” (R. Salakhutdinov and Hinton 2009, 450). The stacking is achieved by treating a hidden layer of one RBM as the visible layer in a second RBM, and so on, until the desired multi-layer architecture is created.

2.2 Degeneracy, instability, and uninterpretability... Oh my!

The highly flexible nature of a RBM (having as it does $n_H + n_V + n_H * n_V$ parameters) creates at least three kinds of potential model impropriety that we will call *degeneracy*, *instability*, and *uninterpretability*. In this section we define these characteristics, consider how to detect them in a RBM, and point out relationships among them.

2.2.1 Near-degeneracy

In Random Graph Model theory, *model degeneracy* means there is a disproportionate amount of probability placed on only a few elements of the sample space, \mathcal{X} , by the model (Handcock 2003). For random graph models, \mathcal{X} denotes all possible graphs that can be constructed from a set of nodes and an exponentially parameterized random graph model has a distribution of the form

$$f_{\boldsymbol{\theta}}(\mathbf{x}) = \frac{\exp(\boldsymbol{\theta}^T \mathbf{t}(\mathbf{x}))}{\gamma(\boldsymbol{\theta})}, \mathbf{x} \in \mathcal{X},$$

where $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^q$ is the model parameter, and $\mathbf{t} : \mathcal{X} \rightarrow \mathbb{R}^q$ is a vector of statistics based on the adjacency matrix of a graph. Here, as earlier, $\gamma(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \mathcal{X}} \exp(\boldsymbol{\theta}^T \mathbf{t}(\mathbf{x}))$ is the normalizing function. Let C denote the convex hull of the potential outcomes of sufficient statistics, $\{\mathbf{t}(\mathbf{x}) : \mathbf{x} \in \mathcal{X}\}$, under the model above. Handcock (2003) classify an exponentially parametrized random graph model at $\boldsymbol{\theta}$ as *near-degenerate* if the mean value of the vector of sufficient statistics under $\boldsymbol{\theta}$, $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \mathbf{t}(\mathbf{X})$, is close to the boundary of C . This makes sense because if a model is near-degenerate in the sense that only a small number of elements of the sample space \mathcal{X} have positive probability, the expected value $\mathbb{E}_{\boldsymbol{\theta}} \mathbf{t}(\mathbf{X})$ is an average of that same small number of values of $\mathbf{t}(\mathbf{x})$ and can be expected to *not* be pulled deep into the interior of C .

A RBM model can be thought to exhibit an analogous form of *near-degeneracy* when there is a disproportionate amount of probability placed on a small number of elements in the sample space of visibles and hiddenes, $\mathcal{C}^{n_V+n_H}$. Using the idea of Handcock (2003), when the random vector $\mathbf{t}(\mathbf{x}) = (v_1, \dots, v_{n_V}, h_1, \dots, h_{n_H}, v_1 h_1, \dots, v_{n_V} h_{n_H}) \in \mathcal{T} \equiv \{\mathbf{t}(\mathbf{x}) : \mathbf{x} \in \mathcal{C}^{n_H+n_V}\}$ from (2), appearing in the neg-potential function $Q_{\boldsymbol{\theta}}(\cdot)$, has a mean vector $\boldsymbol{\mu}(\boldsymbol{\theta}) \in \mathbb{R}^{n_V+n_H+n_V*n_H}$ close to the boundary of the convex hull of \mathcal{T} , then the RBM model can be said to exhibit near-degeneracy at $\boldsymbol{\theta} \in \mathbb{R}^{n_V+n_H+n_H*n_V}$. Here the mean of $\mathbf{t}(\mathbf{x})$ is

$$\begin{aligned} \boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \mathbf{t}(\mathbf{X}) &= \sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \{\mathbf{t}(\mathbf{x}) f_{\boldsymbol{\theta}}(\mathbf{x})\} \\ &= \sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \left\{ \mathbf{t}(\mathbf{x}) \frac{\exp\left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)}{\sum_{\mathbf{x} \in \mathcal{C}^{n_H+n_V}} \exp\left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)} \right\}. \end{aligned}$$

2.2.2 Instability

Considering exponential families of distributions, Schweinberger (2011) introduced a concept of model deficiency related to *instability*. Instability can be roughly thought of as excessive

sensitivity in the model, where small changes in the components of potential data outcomes, \mathbf{x} , may lead to substantial changes in the probability function $f_{\theta}(\mathbf{x})$. To quantify “instability” more rigorously (particularly beyond the definition given by Schweinberger (2011)) it is useful to consider how RBM models might be expanded to incorporate more and more visibles. When increasing the size of RBM models, it becomes necessary to grow the number of model parameters (and in this process one may also arbitrarily expand the number of hidden variables used). To this end, let $\theta_{n_V} \equiv (\theta_{v_1}, \dots, \theta_{v_{n_V}}, \theta_{h_1}, \dots, \theta_{h_{n_H}}, \theta_{11}, \dots, \theta_{n_V n_H})$, $n_V \geq 1$, denote an element of a sequence of RBM parameters indexed by the number n_V of visibles (V_1, \dots, V_{n_V}) and define a (scaled) extremal log-probability ratio of the RBM model at θ_{n_V} as

$$\frac{1}{n_V} \log \left[\frac{\max_{(v_1, \dots, v_{n_V}) \in \mathcal{C}^{n_V}} P_{\theta_{n_V}}(v_1, \dots, v_{n_V})}{\min_{(v_1, \dots, v_{n_V}) \in \mathcal{C}^{n_V}} P_{\theta_{n_V}}(v_1, \dots, v_{n_V})} \right] \equiv \frac{1}{n_V} \text{ELPR}(\theta_{n_V}) \quad (3)$$

where $P_{\theta_{n_V}}(v_1, \dots, v_{n_V}) \propto \sum_{\{h_j\} \in \mathcal{C}} \exp \left(\sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j + \sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j \right)$ is the RBM probability of observing outcome (v_1, \dots, v_{n_V}) for the visible variables (V_1, \dots, V_{n_V}) under parameter vector θ_{n_V} .

In formulating a RBM model for a potentially large number of visibles (i.e., as $n_V \rightarrow \infty$), we will say that the ratio (3) needs to stay bounded for the sequence of RBM models to be stable. That is, we make the following convention.

Definition 1 (S-unstable RBM). Let $\theta_{n_V} \in \mathbb{R}^{n_V + n_H + n_H * n_V}$, $n_V \geq 1$, be an element of a sequence of RBM parameters where the number of hiddens, $n_H \equiv n_H(n_V) \geq 1$, can be an arbitrary function of the number n_V of visibles. A RBM model formulation is *Schweinberger-unstable* or *S-unstable* if

$$\lim_{n_V \rightarrow \infty} \frac{1}{n_V} \text{ELPR}(\theta_{n_V}) = \infty.$$

In other words, given any $c > 0$, there exists an integer $n_c > 0$ so that $\frac{1}{n_V} \text{ELPR}(\theta_{n_V}) > c$ for all $n_V \geq n_c$.

This definition of *S-unstable* is a generalization or re-interpretation of the “unstable” concept of Schweinberger (2011) in that here RBM models for visibles (v_1, \dots, v_{n_V}) do not form an exponential family and the dimensionality of $\boldsymbol{\theta}_{n_V}$ is not fixed, but rather grows with n_V .

S-unstable RBM model sequences are undesirable for several reasons. One is that, as mentioned above, small changes in data can lead to overly-sensitive changes in probability. Consider, for example,

$$\Delta(\boldsymbol{\theta}_{n_V}) \equiv \max \left\{ \log \frac{P_{\boldsymbol{\theta}_{n_V}}(\mathbf{v})}{P_{\boldsymbol{\theta}_{n_V}}(\mathbf{v}^*)} : \mathbf{v} \text{ \& } \mathbf{v}^* \in \mathcal{C}^{n_V} \text{ differ in exactly one component} \right\},$$

denoting the biggest log-probability ratio for a one component change in data outcomes (visibles) at a RBM parameter $\boldsymbol{\theta}_{n_V}$. We then have the following result.

Proposition 1. *Let $c > 0$ and let $ELPR(\boldsymbol{\theta}_{n_V})$ be as in (3) for an integer $n_V \geq 1$. If $\frac{1}{n_V}ELPR(\boldsymbol{\theta}_{n_V}) > c$, then $\Delta(\boldsymbol{\theta}_{n_V}) > c$.*

In other words, if the probability ratio (3) is too large, then a RBM model sequence will exhibit large probability shifts for very small changes in data configurations (i.e., will exhibit instability). Recall the applied example of RBM models as a means to classify images. For data as pixels in an image, the instability result in Proposition 1 manifests itself as a one pixel change in an image (one component of the visible vector) resulting in a large shift in the probability, which in turn could result in a vastly different classification of the image. Examples of this behavior have been presented in Szegedy et al. (2013) for other deep learning models, in which a one pixel change in a test image results in a wildly different classification.

Additionally, S-unstable RBM model sequences are connected to the near-degeneracy of Section 2.2.1 (in which model sequences place all probability on a small portion of their sample spaces). To see this, define an arbitrary modal set of possible outcomes (i.e. set of highest probability outcomes) in RBM models with parameters $\boldsymbol{\theta}_{n_V}, n_V \geq 1$ as

$$M_{\epsilon, \boldsymbol{\theta}_{n_V}} \equiv \left\{ \mathbf{v} \in \mathcal{C}^{n_V} : \log P_{\boldsymbol{\theta}_{n_V}}(\mathbf{v}) > (1 - \epsilon) \max_{\mathbf{v}^*} P_{\boldsymbol{\theta}_{n_V}}(\mathbf{v}^*) + \epsilon \min_{\mathbf{v}^*} P_{\boldsymbol{\theta}_{n_V}}(\mathbf{v}^*) \right\}$$

for a given $0 < \epsilon < 1$. Then S-unstable model sequences are guaranteed to be degenerate, as the following result shows.

Proposition 2. *For an S-unstable RBM model sequence and any $0 < \epsilon < 1$, $P_{\theta_{n_V}}((v_1, \dots, v_{n_V}) \in M_{\epsilon, \theta_{n_V}}) \rightarrow 1$ as $n_V \rightarrow \infty$.*

In other words, S-unstable RBM model sequences are guaranteed to stack up all probability on an arbitrarily narrow set of outcomes for visibles. Proofs of Propositions 1 and 2 appear in Kaplan, Nordman, and Vardeman (2016). These findings also have counterparts in results in Schweinberger (2011), but unlike results there, we do not limit consideration to exponential family forms with a fixed number of parameters.

2.2.3 Uninterpretability

For spatial Markov models, Kaiser (2007) defines a measure of model impropriety he calls *uninterpretability*, which is characterized by dependence parameters in a model being so extreme that marginal mean-structures fail to hold as anticipated by a model statement. We adapt this notion to RBM models as follows. Note that in a RBM, the parameters $\theta_{v_1}, \dots, \theta_{v_{n_V}}$ and $\theta_{h_1}, \dots, \theta_{h_{n_H}}$ are naturally associated with main effects of visible and hidden variables and can be interpreted as (logit functions of) means for variables $V_1, \dots, V_{n_V}, H_1, \dots, H_{n_H}$ in a model with no interaction parameters, $\theta_{ij} = 0, i = 1, \dots, n_V, j = 1, \dots, n_H$. That is, with no interaction parameters, we have from (1) that

$$P_{\theta}(V_i = 1) \propto e^{\theta_{v_i}} \quad \text{and} \quad P_{\theta}(H_j = 1) \propto e^{\theta_{h_j}}, \quad i = 1, \dots, n_V, j = 1, \dots, n_H$$

so that, for example, $\text{logit}(P_{\theta}(V_i = 1)) = \theta_{v_i}$ (or $2\theta_{v_i}$) under the coding $\mathcal{C} = \{0, 1\}$ (or $\{-1, 1\}$). Hence, these main effect parameters have a clear interpretation under an independence model (one with $\theta_{ij} = 0$) but this interpretation can break down as interaction parameters increase in magnitude relative to the size of the main effects. In such cases, the main effect parameters θ_{v_1} and θ_{h_j} are no longer interpretable as originally intended in the models (statements of marginal means) and the dependence parameters are so large as to dominate the entire model probability structure (also destroying the model interpretation of dependence as local conditional modifications of an overall marginal mean structure). To assess which parameter values θ may cause difficulties in interpretation, we use the difference $E[\mathbf{X}|\theta] - E[\mathbf{X}|\theta^*]$ between two model expectations: $E[\mathbf{X}|\theta]$ at θ and

expectations $E[\mathbf{X}|\boldsymbol{\theta}^*]$ where $\boldsymbol{\theta}^*$ matches $\boldsymbol{\theta}$ for all main effects but otherwise has $\theta_{ij} = 0$ for $i = 1, \dots, n_V, j = 1, \dots, n_H$. Hence, $\boldsymbol{\theta}^*$ and $\boldsymbol{\theta}$ have the same main effects but $\boldsymbol{\theta}^*$ has 0 dependence parameters. Uninterpretability is then avoided at a parametric specification $\boldsymbol{\theta}$ if the model expected value at $\boldsymbol{\theta}$ is not very different from the corresponding model expectation under independence. Using this, it is possible to investigate what parametric conditions lead to uninterpretability in a model versus those that guarantee interpretable models. If $E[\mathbf{X}|\boldsymbol{\theta}] - E[\mathbf{X}|\boldsymbol{\theta}^*]$ is large, then the RBM model with parameter vector $\boldsymbol{\theta}$ is said to be uninterpretable. The quantities to compare in the RBM case are

$$E[\mathbf{X}|\boldsymbol{\theta}] = \sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \mathbf{x} f_{\boldsymbol{\theta}}(\mathbf{x}) = \sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \mathbf{x} \frac{\exp\left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)}{\sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \exp\left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)}$$

and

$$E[\mathbf{X}|\boldsymbol{\theta}^*] = \sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \mathbf{x} \frac{\exp\left(\sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)}{\sum_{\mathbf{x} \in \mathcal{C}^{n_V+n_H}} \exp\left(\sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j\right)}$$

3 Explorations

We next explore and numerically explain the relationship between values of $\boldsymbol{\theta}$ and the three notions of model impropriety, near-degeneracy, instability, and uninterpretability, for RBM models of varying sizes.

3.1 Tiny example

To illustrate the ideas of model near-degeneracy, instability, and uninterpretability in a RBM, we consider first the smallest possible (toy) example that consists of one visible node

v_1 and one hidden node h_1 that are both binary. A schematic of this model can be found in Figure 2. Because it seems most common, we shall begin by employing 0/1 encoding of binary variables (both h_1 and v_1 taking values in $\mathcal{C} = \{0, 1\}$). (Eventually we shall argue in Section 3.1.2 that $-1/1$ coding has substantial advantages.)




Figure 2: A small example restricted Boltzmann machine (RBM), with two nodes, one hidden and one visible.

3.1.1 Impropropriety three ways

For this small model, we are able to investigate the symptoms of model impropriety, beginning with near-degeneracy. To this end, recall from Section 2.2.1 that one characterization requires consideration of the convex hull of possible values of statistics $\mathbf{t}(\mathbf{x})$,

$$\mathcal{T} = \{\mathbf{t}(\mathbf{x}) : \mathbf{x} = (v_1, h_1) \in \{0, 1\}^2\} \equiv \{(v_1, h_1, v_1 h_1) : v_1, h_1 \in \{0, 1\}\}$$

appearing in the RBM probabilities for this model. As this set is in three dimensions, we are able to explicitly illustrate the shape of boundary of the convex hull of \mathcal{T} and explore the behavior of the mean vector $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} \mathbf{t}(\mathbf{x})$ as a function of the parameter vector $\boldsymbol{\theta}$. Figure 3 shows the convex hull of our “statistic space,” $\mathcal{T} \subset \{0, 1\}^3$, for this toy problem from two perspectives (enclosed by the unit cube $[0, 1]^3$, the convex hull of $\{0, 1\}^3$). In this small model, note that the convex hull of \mathcal{T} does not fill the unrestricted hull of $\{0, 1\}^3$ because of the relationship between the elements of $\mathcal{T} = \{(v_1, h_1, v_1 h_1) : v_1, h_1 \in \{0, 1\}\}$ (i.e. $v_1 h_1 = 1$ only if $v_1 = h_1 = 1$).

We can compute the mean vector for $\mathbf{t}(\mathbf{x})$ as a function of the model parameters as

$$\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}} [\mathbf{t}(\mathbf{X})] = \sum_{\mathbf{x}=(v_1, h_1) \in \{0, 1\}^2} \left\{ t(\mathbf{x}) \frac{\exp(\theta_{11} h_1 v_1 + \theta_{h1} h_1 + \theta_{v1} v_1)}{\gamma(\boldsymbol{\theta})} \right\} = \begin{bmatrix} \frac{\exp(\theta_{v1}) + \exp(\theta_{11} + \theta_{v1} + \theta_{h1})}{\gamma(\boldsymbol{\theta})} \\ \frac{\exp(\theta_{h1}) + \exp(\theta_{11} + \theta_{v1} + \theta_{h1})}{\gamma(\boldsymbol{\theta})} \\ \frac{\exp(\theta_{11} + \theta_{v1} + \theta_{h1})}{\gamma(\boldsymbol{\theta})} \end{bmatrix}$$




Figure 3: The convex hull of the "statistic space" in three dimensions for the toy RBM with one visible and one hidden node. The convex hull of $\mathcal{T} = \{\mathbf{t}(\mathbf{x}) : \mathbf{x} \in \mathcal{C}^{n_V+n_H}\}$ does not fill the unit cube because of the relationship between the elements of \mathcal{T} .

where $\gamma(\boldsymbol{\theta}) = \sum_{h_1=0}^1 \sum_{v_1=0}^1 \exp(\theta_{11}h_1v_1 + \theta_{h_1}h_1 + \theta_{v_1}v_1)$. The three parametric coordinate functions of $\boldsymbol{\mu}(\boldsymbol{\theta})$ can be represented as in Figure 4. (Contour plots for three coordinate functions are shown in columns for various values of θ_{11} .) In examining these, we see that as coordinates of $\boldsymbol{\theta}$ grow larger in magnitude, at least one mean function for the entries of $\mathbf{t}(\mathbf{x})$ approaches a value 0 or 1, forcing $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}}\mathbf{t}(\mathbf{x})$ to be near to the boundary of the convex hull of \mathcal{T} , as a sign of model near-degeneracy. Thus, for a very small example we can see the relationship between values of $\boldsymbol{\theta}$ and model degeneracy.

Secondly, we can look at $\text{ELPR}(\boldsymbol{\theta})$ from (3) for this tiny model in order to consider model instability as a function of RBM parameters. Recall that large values of $\text{ELPR}(\boldsymbol{\theta})$ are associated with an extreme sensitivity of the model probabilities $f_{\boldsymbol{\theta}}(\mathbf{x})$ to small changes in \mathbf{x} (see Proposition 1). The quantity $\text{ELPR}(\boldsymbol{\theta})$ for this small RBM is

$$\text{ELPR}(\boldsymbol{\theta}) = \log \left[\frac{\max_{(v_1, \dots, v_{n_V}) \in \mathcal{C}^{n_V}} P_{\boldsymbol{\theta}_{n_V}}(v_1, \dots, v_{n_V})}{\min_{(v_1, \dots, v_{n_V}) \in \mathcal{C}^{n_V}} P_{\boldsymbol{\theta}_{n_V}}(v_1, \dots, v_{n_V})} \right] = \log \left[\frac{\max_{v_1 \in \mathcal{C}} \sum_{h_1 \in \mathcal{C}} \exp \{ \theta_{11}h_1v_1 + \theta_{h_1}h_1 + \theta_{v_1}v_1 \}}{\min_{v_1 \in \mathcal{C}} \sum_{h_1 \in \mathcal{C}} \exp \{ \theta_{11}h_1v_1 + \theta_{h_1}h_1 + \theta_{v_1}v_1 \}} \right].$$

Figure 5 shows contour plots of $\text{ELPR}(\boldsymbol{\theta})/n_V$ for various values of $\boldsymbol{\theta}$ in this model with

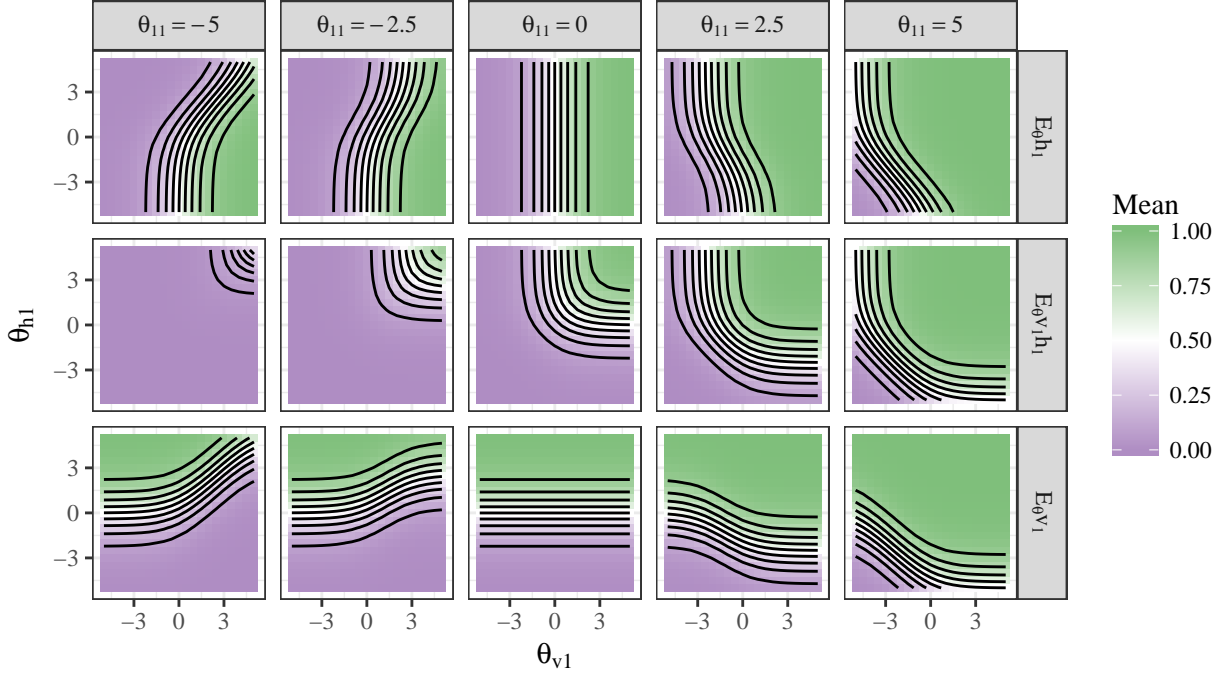


Figure 4: Contour plots for the three parametric mean functions of sufficient statistics for a RBM with one visible and one hidden node.

$n_v = 1$. We can see that this quantity is large for large magnitudes of θ , especially for large values of the dependence/interaction parameter θ_{11} . This suggests instability as $|\theta|$ becomes large, agreeing also with the concerns about near-degeneracy produced by consideration of $\mu(\theta)$.

Finally to consider the effect of θ on potential model uninterpretability, we can look at the difference between model expectations, $E[\mathbf{X}|\theta]$, and expectations given independence, $E[\mathbf{X}|\theta^*]$ for the tiny toy RBM model where $\mathbf{X} = (V_1, H_1, V_1 H_1)$. This difference is given by

$$E[\mathbf{X}|\theta] - E[\mathbf{X}|\theta^*] = \begin{bmatrix} \frac{\exp(\theta_{11} + \theta_{v_1} + 2\theta_{h_1}) - \exp(\theta_{v_1} + 2\theta_{h_1})}{(\exp(\theta_{v_1}) + \exp(\theta_{h_1}) + \exp(\theta_{11} + \theta_{v_1} + \theta_{h_1}))(\exp(\theta_{v_1}) + \exp(\theta_{h_1}) + \exp(\theta_{v_1} + \theta_{h_1}))} \\ \frac{\exp(\theta_{11} + 2\theta_{v_1} + \theta_{h_1}) - \exp(2\theta_{v_1} + \theta_{h_1})}{(\exp(\theta_{v_1}) + \exp(\theta_{h_1}) + \exp(\theta_{11} + \theta_{v_1} + \theta_{h_1}))(\exp(\theta_{v_1}) + \exp(\theta_{h_1}) + \exp(\theta_{v_1} + \theta_{h_1}))} \end{bmatrix}.$$

Again, we can inspect these coordinate functions of this vector difference to look for a relationship between parameter values and large values of $E[\mathbf{X}|\theta] - E[\mathbf{X}|\theta^*]$ as a signal of uninterpretability for the toy RBM.




Figure 5: $\text{ELPR}(\boldsymbol{\theta})/n_v$ for various values of $\boldsymbol{\theta}$ for the tiny example model. Recall here n_v is the number of visible nodes and here is 1. This quantity is large for large magnitudes of $\boldsymbol{\theta}$.




Figure 6: The absolute difference between coordinates of model expectations, $E[\mathbf{X}|\boldsymbol{\theta}]$, and expectations given independence, $E[\mathbf{X}|\boldsymbol{\theta}^*]$ for a RBM with one visible and one hidden node. As an indicator of uninterpretability, note that differences in expectations increase as the dependence parameter θ_{11} deviates from zero.

Figure 6 shows that the absolute difference between coordinates of the vector of model expectations, $E[\mathbf{X}|\boldsymbol{\theta}]$ and corresponding expectations $E[\mathbf{X}|\boldsymbol{\theta}^*]$ given independence grow for the toy RBM as the values of $\boldsymbol{\theta}$ are farther from zero, especially for large magnitudes of the dependence parameter θ_{11} . This is a third indication that parameter vectors of large magnitude lead to model impropriety in a RBM.

3.1.2 Data coding to mitigate apparent degeneracy

Multiple encodings of the binary variables are possible. For example, we could allow hiddens $(H_1, \dots, H_{n_H}) \in \{0, 1\}^{n_H}$ and visibles $(V_1, \dots, V_{n_V}) \in \{0, 1\}^{n_V}$, as in the previous sections or we could instead encode the state of the variables as $\{-1, 1\}^{n_H}$ and $\{-1, 1\}^{n_V}$. This will result in variables $\mathbf{t}(\mathbf{X})$ from (2) satisfying $\mathbf{t}(\mathbf{x}) \in \{0, 1\}^{n_H+n_V+n_H*n_V}$ or $\mathbf{t}(\mathbf{x}) \in \{-1, 1\}^{n_H+n_V+n_H*n_V}$ depending on how we encode “on” and “off” states in the nodes. To explore the effect of this encoding on the potential for apparent near-degeneracy of a RBM, we can consider the ratio of the volume of the hypercube with corners at elements of \mathcal{T} to the volume of the convex hull of \mathcal{T} under both possible encodings (i.e., compare convex hull of \mathcal{T} to either $[0, 1]^3$ or $[-1, 1]^3$ under the respective encodings).

For the small two node example, the 0/1 encoding loses 83.33% of the volume of $[0, 1]^3$ and the $-1/1$ encoding loses 66.67% of the volume of $[-1, 1]^3$. While there is clearly a one-to-one mapping between models for these two possible codings, this notion of lost volume can be helpful to geometrically conceptualize how difficult it will be for the mean vector $\boldsymbol{\mu}(\boldsymbol{\theta})$ to avoid the boundary of \mathcal{T} , and thus avoid apparent near-degeneracy. In fact, if we look at the ratio of volume within the convex hull defined by \mathcal{T} and the corresponding hypercube, we can see a relationship emerge as the number of nodes (and thus parameters) increases. In Figure 7, it is evident that as the number of parameters increases, this ratio is decreasing at an increasing rate, meaning it gets more and more difficult to avoid the boundary of the convex hull and thus appearance of near degeneracy. Additionally, it appears that the $-1/1$ encoding suffers slightly less from this problem. This suggests that if intuition about models is to be formed in terms of a data encoding that encourages of “non-pathological” means $\boldsymbol{\mu}(\boldsymbol{\theta}) = E_{\boldsymbol{\theta}}(\mathbf{t}(\mathbf{x}))$ for $\mathbf{t}(\mathbf{x})$, then the $-1/1$ encoding is preferable to 0/1 coding.




Figure 7: The relationship between volume of the convex hull of possible values of the RBM sufficient statistics and the cube containing it for different configurations of nodes.

In addition to the argument that the $-1/1$ data encoding mitigates some perceived prevalence of near-degeneracy it also has the benefit of providing a guaranteed-to-be non-degenerate model at $\boldsymbol{\theta} = \mathbf{0} \in \mathbb{R}^{n_H+n_V+n_H*n_V}$, where the zero vector then serves as the natural center of the parameter space and induces the simplest possible model properties for the RBM (i.e., at $\boldsymbol{\theta} = \mathbf{0}$, all variables are independent and visible variables are independent and uniformly distributed on $\{-1, 1\}^{n_V}$). The proof of this and further exploration of the equivalence of the $\boldsymbol{\theta}$ parameterization of the RBM model class and parameterization by $\boldsymbol{\mu}(\boldsymbol{\theta})$ is in the on-line supplementary materials. Hence, while from some computing perspectives 0/1 coding might seem most natural, the $-1/1$ coding is far more convenient and interpretable from the point of view of statistical modeling, where it makes parameters simply interpreted in terms of symmetrically defined main effects and interactions. In light of all of these matters we will henceforth employ the $-1/1$ coding.

3.2 Exploring manageable examples

To explore the impact of RBM parameter vector $\boldsymbol{\theta}$ magnitude on near-degeneracy, instability, and uninterpretability, we consider models of small size. For $n_H, n_V \in \{1, \dots, 4\}$, we sample

100 values of $\boldsymbol{\theta}$ with various magnitudes (details to follow). For each set of parameters we then calculate metrics of model impropriety introduced in Section 2.2 based on $\boldsymbol{\mu}(\boldsymbol{\theta})$, $\text{ELPR}(\boldsymbol{\theta})/n_V$, and the absolute coordinates of $\text{E}[\mathbf{X}|\boldsymbol{\theta}] - \text{E}[\mathbf{X}|\boldsymbol{\theta}^*]$. In the case of near-degeneracy, we classify each model as near-degenerate or “viable” based on the distance of $\boldsymbol{\mu}(\boldsymbol{\theta})$ from the boundary of the convex hull of \mathcal{T} and look at the fraction of models that are “near-degenerate,” meaning they are within a small distance $\epsilon > 0$ of the boundary of the convex hull. We define “small” through a rough estimation of the volume of the hull for each model size. We pick $\epsilon_0 = 0.05$ for $n_H = n_V = 1$ and then, for every other n_H and n_V , set $m = n_H + n_V + n_V * n_H$ and pick ϵ so that $1 - (1 - 2\epsilon_0)^3 = 1 - (1 - 2\epsilon)^m$. In this way, we roughly scale the volume of the “small distance” to the boundary of the convex hull to be equivalent across model dimensions.

In our numerical experiment, we split $\boldsymbol{\theta} = (\boldsymbol{\theta}_{main}, \boldsymbol{\theta}_{interaction})$ into $\boldsymbol{\theta}_{main}$ and $\boldsymbol{\theta}_{interaction}$, in reference to which variables in the probability function the parameters correspond (whether they multiply a v_i or a h_j or they multiply a $v_i h_j$), and allow the two types of terms to have varying average magnitudes, $\|\boldsymbol{\theta}_{main}\|/(n_H + n_V)$ and $\|\boldsymbol{\theta}_{interaction}\|/(n_H * n_V)$. These average magnitudes vary on a grid between 0.001 and 3 with 24 breaks, yielding 576 grid points. (By looking at the average magnitudes, we are able to later consider the potential benefit of shrinking each parameter value θ_i towards zero in a Bayesian fitting technique.) At each point in the grid, 100 vectors ($\boldsymbol{\theta}_{main}$) are sampled uniformly on a sphere with radius corresponding to the first coordinate in the grid and 100 vectors ($\boldsymbol{\theta}_{interaction}$) are sampled uniformly on a sphere with radius corresponding to the second coordinate in the grid via sums of squared and scaled iid Normal(0, 1) variables. These vectors are then paired to create 100 values of $\boldsymbol{\theta}$ with magnitudes at each point in the grid.

The results of this numerical study are summarized in Figures 8, 9, and 10. From these three figures, it is clear that all three measures of model impropriety show higher values for larger magnitudes of the parameter vectors. Additionally, since there are $n_H * n_V$ interaction terms in $\boldsymbol{\theta}$ versus only $n_H + n_V$ main effect terms, for large models there are many more interaction parameters than main effects in the models. And so, severely limiting the magnitude of the individual interactions may well help prevent model impropriety.

Figure 11 shows the fraction of near-degenerate models for each magnitude of $\boldsymbol{\theta}$ for each

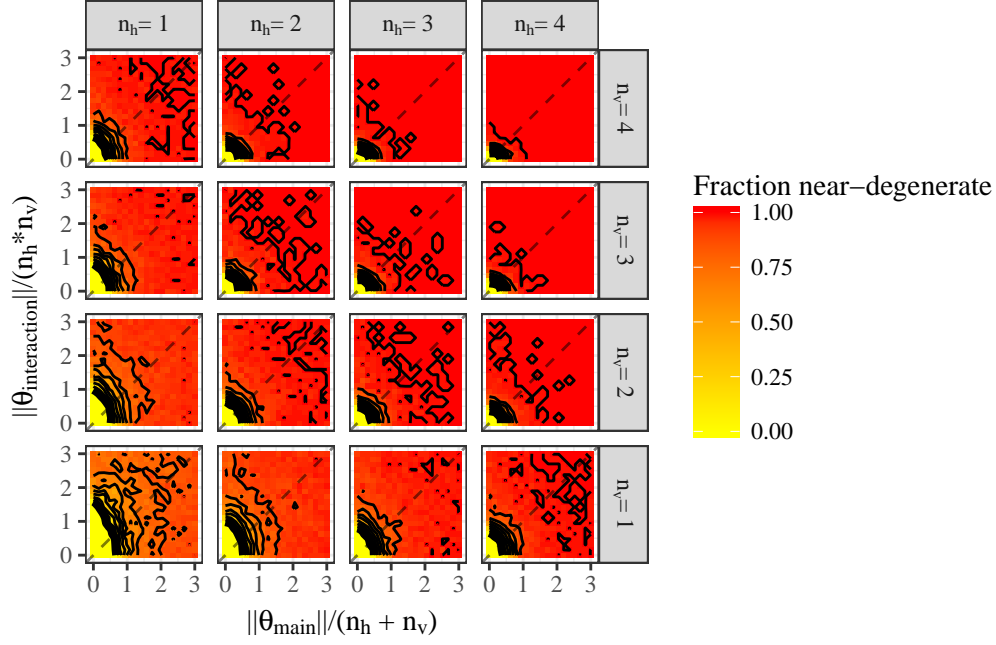


Figure 8: Results from the numerical experiment, here looking at the fraction of models that were near-degenerate for each combination of magnitude of θ and model size. Black lines show the contour levels for fraction of near-degeneracy, while the thick black line shows the level where the fraction of near-degenerate models is .05.

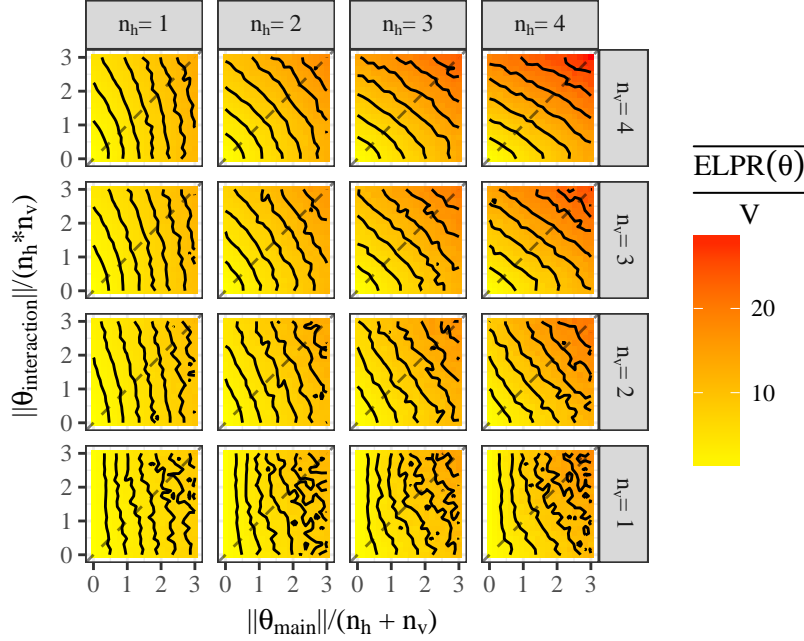


Figure 9: Results from the experiment, here looking at the sample mean value of $\text{ELPR}(\theta)/V$ at each grid point for each combination of magnitude of θ and model size. As the magnitude of θ grows, so does the value of this metric, indicating typical instability in the model.

model architecture. For each number n_v of visibles in the model, as the number n_h of hiddens increase, the fraction near-degenerate diverges from zero at increasing rates for larger values of $\|\theta\|$. This shows that as model size gets larger, the risk of degeneracy starts at a slightly larger magnitude of parameters, but very quickly increases until reaching close to 1.

These manageable examples indicate that RBMs are near-degenerate, unstable, and uninterpretable for large portions of the parameter space with large $\|\theta\|$. This, however, is not the only potential problem to be faced when using these models. There is the matter of principled/rigorous fitting of RBM models.

4 Model Fitting

Typically, fitting a RBM via maximum likelihood (ML) methods will be infeasible due mainly to the intractability of the normalizing term $\gamma(\theta)$ in a model (1) of any realistic size.

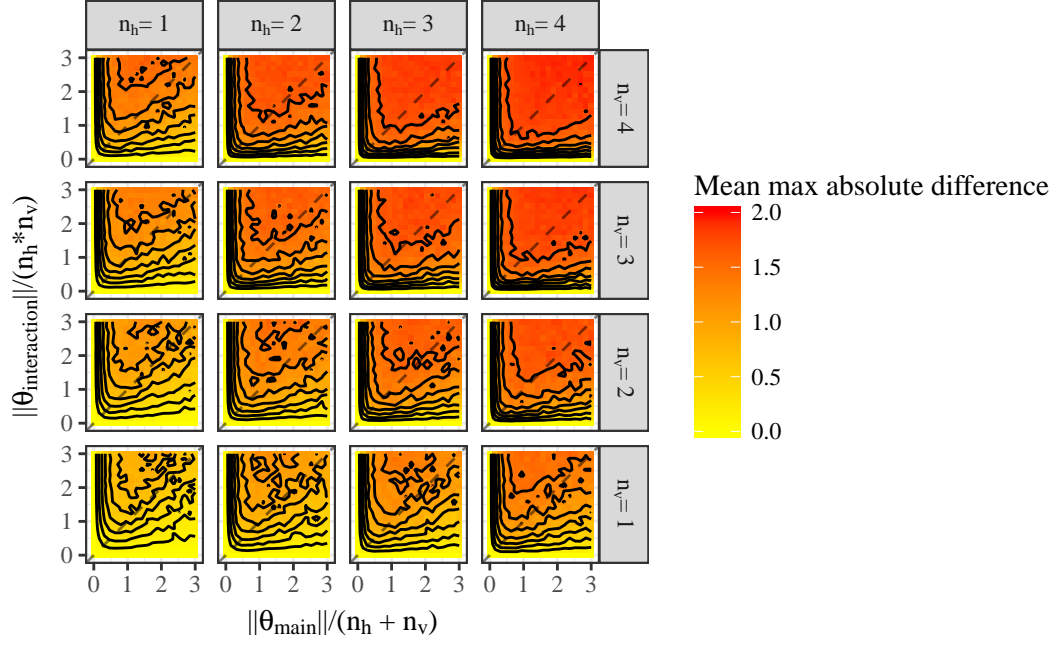


Figure 10: The sample mean of the maximum component of the absolute difference between the model expectation vector, $E[\mathbf{X}|\boldsymbol{\theta}]$, and the expectation vector given independence, $E[\mathbf{X}|\boldsymbol{\theta}^*]$. Larger magnitudes of $\boldsymbol{\theta}$ correspond to larger differences, thus indicating reduced interpretability.

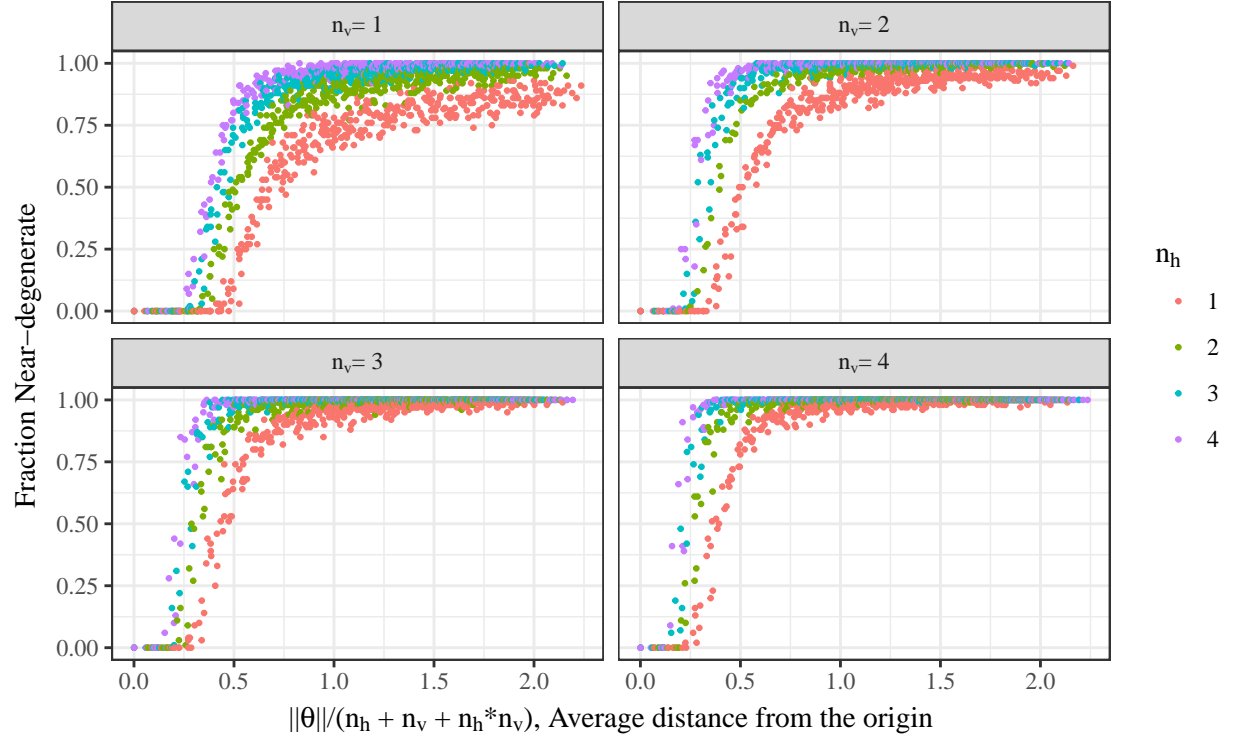


Figure 11: The fraction of near-degenerate models for each magnitude of θ . For each number v of visibles in the model, the fraction near-degenerate becomes greater than zero at larger values of $\|\theta\|$ as n_H increases and the slope becomes steeper as n_H increases as well.

Ad hoc methods are used instead, which aim to avoid this problem by using stochastic ML approximations that employ a small number of MCMC draws (i.e., contrastive divergence, (Hinton 2002)).

However, computational concerns are not the only issues with fitting a RBM using ML. In addition, a RBM model, with the appropriate choice of parameters and number of hidden, has the potential to re-create any distribution for the data (i.e., reproduce any specification of cell probabilities for the binary data outcomes). For example, Montufar and Ay (2011) show that any distribution on $\{0, 1\}^{n_v}$ can be approximated arbitrarily well by a RBM with $2^{n_v-1} - 1$ hidden units. We provide a small example that illustrates that in fact there can be many such approximations.

For simplicity, consider a model with two visible variables (V_1, V_2) and one hidden H_1 . In this case, there are four possible data realizations for (V_1, V_2) given by $(\pm 1, \pm 1)$ and we may express the model probabilities as

$$P(V_1 = v_1, V_2 = v_2 | \theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \theta_{11}, \theta_{21}) \propto \exp(v_1 \theta_{v_1} + v_2 \theta_{v_2}) \sum_{h \in \{\pm 1\}} \exp(h[\theta_{h_1} + \theta_{11} v_1 + \theta_{21} v_2]),$$

for $(v_1, v_2) \in \{-1, 1\}^2$, in terms of real-valued parameters $\theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \theta_{11}, \theta_{21}$. Given any specified cell probabilities, say

$$0 \leq p_{(-1,-1)}, p_{(1,-1)}, p_{(-1,1)}, p_{(1,1)}, \quad (4)$$

for the outcomes $(\pm 1, \pm 1)$, values for parameters $(\theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \theta_{11}, \theta_{21})$ may be chosen to approximate such cell probabilities with arbitrary closeness. In fact, when the cell probabilities (4) are all strictly positive, parameters in the RBM model can be specified to reflect these probabilities exactly. And, when one or more of the cell probabilities (4) are zero, the corresponding RBM probabilities $P(V_1 = v_1, V_2 = v_2 | \theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \theta_{11}, \theta_{21})$ may never be identically zero (due to exponential terms in the model) but parameters can be still selected to make the appropriate RBM cell probabilities arbitrarily small.

To demonstrate, we assume $p_{(-1,-1)} > 0$ (without loss of generality) in the specified cell probabilities (4) and replace parameters θ_{11}, θ_{21} with $\Delta_1 \equiv \theta_{11} + \theta_{21}$ and $\Delta_2 \equiv \theta_{11} - \theta_{21}$.

We may then prescribe values of $\theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \Delta_1, \Delta_2$ so that the model probability ratio

$$P(V_1 = v_1, V_2 = v_2 | \theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \Delta_1, \Delta_2) / P(V_1 = -1, V_2 = -1 | \theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \Delta_1, \Delta_2)$$

matches the corresponding ratio $p_{(v_1, v_2)} / p_{(-1, -1)}$ over three values of $(v_1, v_2) = (1, -1), (-1, 1), (1, 1)$.

For instance, assuming the cell probabilities from (4) are all positive, these probabilities can be exactly reproduced by choosing

$$\begin{aligned}\theta_{v_1} &= \frac{1}{2} \log \left(\frac{p_{(1, -1)} \exp(\theta_{h_1} - \Delta_1) + \exp(-\theta_{h_1} + \Delta_1)}{p_{(-1, -1)} \exp(\theta_{h_1} + \Delta_2) + \exp(-\theta_{h_1} - \Delta_2)} \right), \\ \theta_{v_2} &= \frac{1}{2} \log \left(\frac{p_{(-1, 1)} \exp(\theta_{h_1} - \Delta_1) + \exp(-\theta_{h_1} + \Delta_1)}{p_{(-1, -1)} \exp(\theta_{h_1} - \Delta_2) + \exp(-\theta_{h_1} + \Delta_2)} \right)\end{aligned}$$

and selecting $\theta_{h_1}, \Delta_1, \Delta_2$ to solve

$$\frac{p_{(1, 1)} p_{(-1, -1)}}{p_{(-1, 1)} p_{(1, -1)}} = \frac{\ell(|\theta_{h_1}|) + \ell(|\Delta_1|)}{\ell(|\theta_{h_1}|) + \ell(|\Delta_2|)}, \quad (5)$$

based on a monotonically increasing function $\ell(x) \equiv \exp(-2x) + \exp(2x)$, $x \geq 0$. If $[p_{(1, 1)} p_{(-1, -1)}] / [p_{(-1, 1)} p_{(1, -1)}] \geq 1$, one can pick any values for $\theta_{h_1}, \Delta_2 \in \mathbb{R}$ and solve (5) for $|\Delta_1|$; likewise, when $[p_{(1, 1)} p_{(-1, -1)}] / [p_{(-1, 1)} p_{(1, -1)}] < 1$ in (5), one may solve for $|\Delta_2|$ upon choosing any values for $\theta_{h_1}, \Delta_1 \in \mathbb{R}$.

Alternatively, if exactly one specified cell probability in (4) is zero, say $p_{(1, 1)}$ (without loss of generality), we can select parameters $\theta_{v_1}, \theta_{v_2}$ as above based on a sequence $(\theta_{h_1}, \Delta_1, \Delta_2) \equiv (\theta_{h_1}^{(m)}, \Delta_1^{(m)}, \Delta_2^{(m)})$, $m \in \{1, 2, \dots\}$ of the remaining parameter values such that $\lim_{m \rightarrow \infty} |\Delta_1^{(m)}| = \infty$ and $\lim_{m \rightarrow \infty} (|\theta_{h_1}^{(m)}| + |\Delta_2^{(m)}|) / |\Delta_1^{(m)}| = 0$ hold. This guarantees that the resulting RBM model matches the given cell probabilities (4) in the limit:

$$\lim_{m \rightarrow \infty} P(V_1 = v_1, V_2 = v_2 | \theta_{v_1}, \theta_{v_2}, \theta_{h_1}, \Delta_1, \Delta_2) = p_{(v_1, v_2)}, \quad (v_1, v_2) \in \{(\pm 1, \pm 1)\}. \quad (6)$$

If exactly two specified probabilities in (4) are zero, say $p_{(1, 1)}$ and $p_{(-1, 1)}$ (without loss of generality), then a limit approximation as in (6) follows by picking θ_{v_1} as above based on any choices of $(\theta_{h_1}, \Delta_1, \Delta_2)$ and choosing a sequence of $\theta_{v_2} \equiv \theta_{v_2}^{(m)}$ values for which $\theta_{v_2}^{(m)} \rightarrow -\infty$.

The previous discussion illustrates the fact that the RBM model class suffers from parameter identifiability issues that go beyond mere symmetries in the parametrization. Not only it is possible to approximate any distribution on the visibles arbitrarily well (cf. Montufar and Ay 2011), but quite different parameter settings can induce the same essential RBM model. However, this is not the most disastrous implication of the RBM parameterization. A far worse consequence is that, when fitting the RBM model by likelihood-based methods, we already know the nature of the answer before we begin: namely, such fitting will simply aim to reproduce the empirical distribution from the training data if sufficiently many hidden units are in the model. That is, based on a random sample of vectors of visible variables, the model for the cell probabilities that has the highest likelihood over *all possible model classes* (i.e., RBM-based or not) is the empirical distribution, and the over-parametrization of the RBM model itself ensures that this empirical distribution can be arbitrarily well-approximated.

For illustration, continue the simple example from above with n iid observations, each consisting of two realized visibles (V_1, V_2) . In which case, when the specified cell probabilities $p_{(-1,-1)}, p_{(1,-1)}, p_{(-1,1)}, p_{(1,1)}$ in (4) are taken as the empirical cell frequencies from the sample, there is no better model based on maximum likelihood, and the discussion above (cf. (6)) shows that RBM model parameters can be chosen to re-create this empirical distribution to an arbitrarily close degree. Hence, RBM model fitting based on ML will simply seek to reproduce the empirical distribution. What's more, whenever this empirical distribution contains empty cells, fitting steps for the RBM model will necessarily aim to choose parameters that necessarily diverge to infinity in magnitude in order to zero-out the corresponding RBM cell probabilities. In data applications with a large sample space, it is unlikely that the training set will include at least one of each possible vector outcome (unlike this small example). This implies that some RBM model parameters must diverge to $+\infty$ to mimic the empirical distribution with empty cells and, as we have already discussed in Section 3, large magnitudes of θ lead to model impropriety in the RBM.

Here we consider what might be done in a principled manner to prevent both overfitting and model impropriety, testing on a $n_V = n_H = 4$ case that already stretches the limits of what is computable - in particular we consider Bayes methods.

4.1 Bayesian model fitting

To avoid model impropriety for a fitted RBM, we want to avoid parts of the parameter space $\mathbb{R}^{n_V+n_H+n_V*n_H}$ that lead to near-degeneracy, instability, and uninterpretability. Motivated by the insights in Section 3.2, one idea is to shrink $\boldsymbol{\theta} = (\boldsymbol{\theta}_{main}, \boldsymbol{\theta}_{interaction})$ toward $\mathbf{0}$ by specifying priors that place low probability on large values of $\|\boldsymbol{\theta}\|$, specifically focusing on shrinking $\boldsymbol{\theta}_{interaction}$ more than $\boldsymbol{\theta}_{main}$. This is similar to an idea advocated by Hinton (2010) called *weight decay*, in which a penalty is added to the interaction terms in the model, $\boldsymbol{\theta}_{interaction}$, shrinking their magnitudes.

Table 1: Parameters used to fit a test case with $V = H = 4$. This parameter vector was chosen as a sampled value of $\boldsymbol{\theta}$ that was not near the convex hull of the sufficient statistics for a grid point in Figure 8 with $< 5\%$ near-degeneracy.

Parameter	Value	Parameter	Value	Parameter	Value
θ_{v1}	-1.104376	θ_{11}	-0.0006334	θ_{31}	-0.0038301
θ_{v2}	-0.2630044	θ_{12}	-0.0021401	θ_{32}	0.0032237
θ_{v3}	0.3411915	θ_{13}	0.0047799	θ_{33}	0.0020681
θ_{v4}	-0.2583769	θ_{14}	0.0025282	θ_{34}	0.0041429
θ_{h1}	-0.1939302	θ_{21}	0.0012975	θ_{41}	0.0089533
θ_{h2}	-0.0572858	θ_{22}	0.0000253	θ_{42}	-0.0042403
θ_{h3}	-0.2101802	θ_{23}	-0.0004352	θ_{43}	-0.000048
θ_{h4}	0.2402456	θ_{24}	-0.0086621	θ_{44}	0.0004767

We considered a test case with $n_V = n_H = 4$ and parameters given in Table 1. This parameter vector was chosen as a sampled value of $\boldsymbol{\theta}$ that was not near the convex hull of the space of values of sufficient statistics for a grid point in Figure 8 with $< 5\%$ near-degeneracy. We simulated $n = 5,000$ realizations of visibles as a training set and fit the RBM using three Bayes methodologies. These involved the following:

1. A “*trick*” prior. Here we cancel out normalizing term in the likelihood so that resulting full conditionals of $\boldsymbol{\theta}$ are multivariate Normal. The h_j are carried along in the MCMC

sampling from the posterior as latent variables.

$$\pi(\boldsymbol{\theta}) \propto \gamma(\boldsymbol{\theta})^n \exp \left(-\frac{1}{2C_1} \boldsymbol{\theta}'_{main} \boldsymbol{\theta}_{main} - \frac{1}{2C_2} \boldsymbol{\theta}'_{interaction} \boldsymbol{\theta}_{interaction} \right),$$

where

$$\gamma(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \mathcal{C}^{n_H+n_V}} \exp \left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^V \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j \right) \text{ and } C_2 < C_1$$

This is the method of Li (2014). We will refer to this method as Bayes with Trick Prior and Latent Variables (BwTPLV).

2. *A truncated Normal prior.* Here we use independent spherical normal distributions as priors for $\boldsymbol{\theta}_{main}$ and $\boldsymbol{\theta}_{interaction}$, with $\sigma_{interaction} < \sigma_{main}$, *truncated* at $3\sigma_{main}$ and $3\sigma_{interaction}$, respectively. Full conditional distributions are not conjugate, and simulation from the posterior was accomplished using a geometric adaptive Metropolis Hastings step (Zhou 2014) and calculation of likelihood normalizing constant. (This computation is barely feasible for a problem of this size and would be infeasible for larger problems.) Here the h_j are carried along in the MCMC implementation as latent variables. We will refer to this method as Bayes with Truncated Normal prior and Latent Variables (BwTNLV).
3. *A truncated Normal prior and marginalized likelihood.* Here we marginalize out \mathbf{h} in $f_{\boldsymbol{\theta}}(\mathbf{x})$, and use the truncated Normal priors applied to the marginal probabilities for visible variables given by

$$g_{\boldsymbol{\theta}}(\mathbf{v}) \propto \sum_{\mathbf{h} \in \mathcal{C}^{n_H}} \exp \left(\sum_{i=1}^{n_V} \sum_{j=1}^{n_H} \theta_{ij} v_i h_j + \sum_{i=1}^{n_V} \theta_{v_i} v_i + \sum_{j=1}^{n_H} \theta_{h_j} h_j \right), \mathbf{v} \in \mathcal{C}^{n_V}.$$

We will refer to this method as Bayes with Truncated Normal prior and Marginalized Likelihood (BwTNML).

The three fitting methods are ordered by computational feasibility in a real-data situation, with BwTPLV being the most computationally feasible due to conjugacy and BwTNML the least feasible due to the marginalization and need for an adaptive Metropolis Hastings

step. All three methods require choosing the values of hyperparameters. In each case, we have chosen these values based on a rule of thumb that shrinks $\theta_{interaction}$ more than θ_{main} . Additionally, BwTPLV requires additional tuning to choose C_1 and C_2 , reducing its appeal. The values used for the hyperparameters in our simulation are presented in Table 2.

Table 2: The values used for the hyperparameters for all three fitting methods. A rule of thumb is imposed which decreases prior variances for the model parameters as the size of the model increases and also shrinks $\theta_{interaction}$ more than θ_{main} . The common C defining C_1 and C_2 in the BwTPLV method is chosen by tuning.

Method	Hyperparameter	Value
BwTPLV	C_1	$\frac{C}{n} \frac{1}{n_H + n_V}$
	C_2	$\frac{C}{n} \frac{1}{n_H * n_V}$
BwTNLV	σ_{main}^2	$\frac{1}{n_H + n_V}$
	$\sigma_{interaction}^2$	$\frac{1}{n_H * n_V}$
BwTNML	σ_{main}^2	$\frac{1}{n_H + n_V}$
	$\sigma_{interaction}^2$	$\frac{1}{n_H * n_V}$

It should be noted that BwTNLV (2.) and BwTNML (3.) are drawing from the same stationary posterior distribution for vectors of visibles. The difference between these two methods is in how well the chains mix and how quickly they arrive at the target posterior distribution. After a burn-in period of 50 iterations selected by inspecting the trace plots, we assess the issue of mixing in two ways. First, the autocorrelation functions (ACF) from each posterior sample corresponding to a model probability for a visible vector outcome $\mathbf{v} = (v_1, v_2, v_3, v_4) \in \{\pm 1\}^4$ (i.e., computed from θ under (1)) are assessed and plotted in Figure 12 with BwTNLV in black and BwTNML in red. As expected, ACF corresponding to the method that marginalizes out the hidden variables from the likelihood decreases to zero at a much faster rate, indicating better mixing for the chain.

Secondly, we can assess the mixing of our chains using an idea of effective sample size. If the MCMC chain were truly iid draws from the target distribution, then for the parameter $p^{(i)}$ denoting the probability of the i th vector outcome for the four visibles $\mathbf{v} = (v_1, v_2, v_3, v_4) \in \{\pm 1\}^4$, $i = 1, \dots, 16$, its estimate as the average $\bar{p}^{(i)}$ of posterior sample versions would

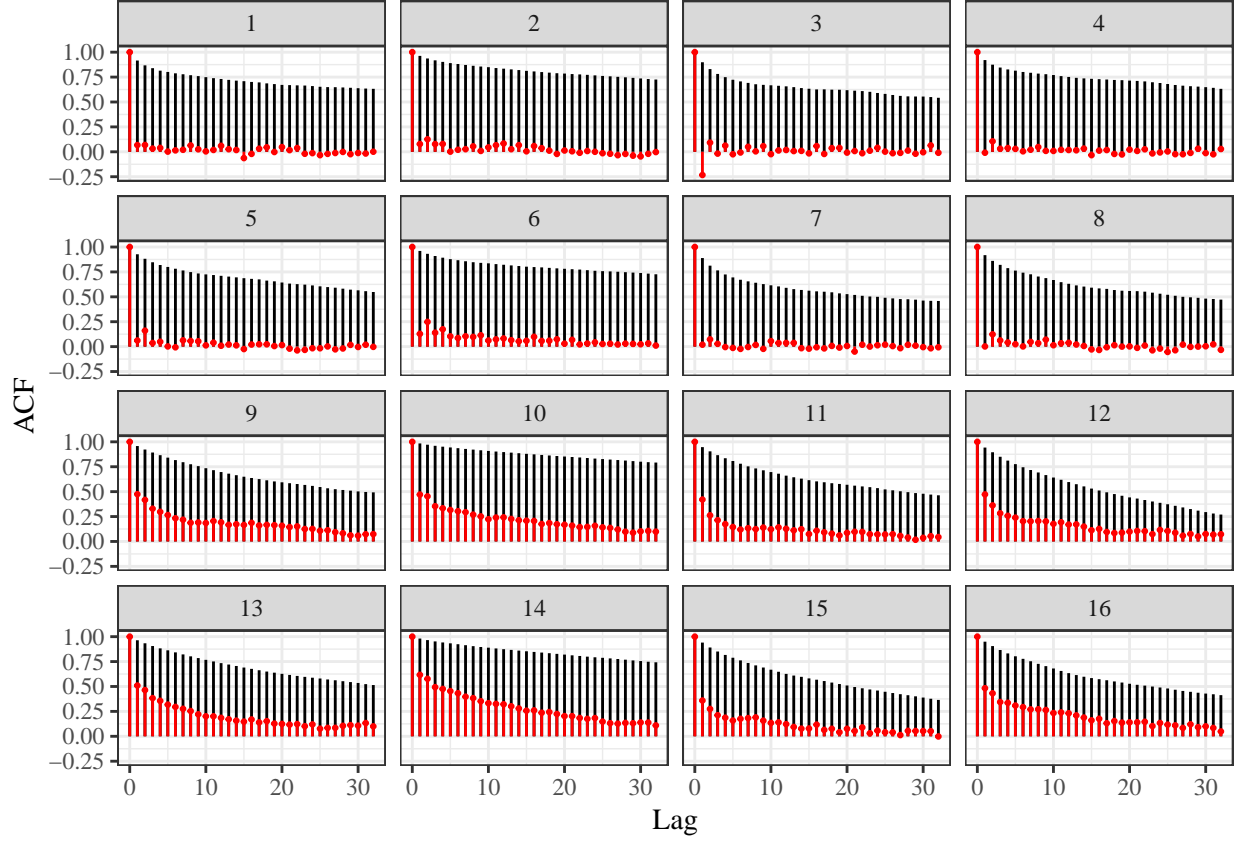


Figure 12: The autocorrelation functions (ACF) for the posterior probabilities of all $2^4 = 16$ possible outcomes for the vector of four visibles assessed at multiple lags for each method with BwTNLV in black and BwTNML in red. As expected, ACF corresponding to the method that marginalizes out the hidden variables from the likelihood decreases to zero at a much faster rate, indicating better mixing for the chain.

be approximately Normal with mean the true posterior marginal probability of $p^{(i)}$, and variance σ_i^2/M , where σ_i^2 is the true posterior variance of $p^{(i)}$ and M is the length of the chain. However, with the presence of correlation in our chain, the asymptotic variance of $\bar{p}^{(i)}$ is instead approximately some C_i/M , where C_i is some positive constant such that $C_i > \sigma_i^2$. We can use an overlapping block-means approach (Gelman, Shirley, and others 2011) to get a crude estimate for C_i as $\hat{C}_i = bS_b^2$, where S_b^2 denotes the sample variance of overlapping block means $\{\bar{p}_j^{(i)} = \sum_{k=j}^{j+b-1} p_k^{(i)} / b\}_{j=1}^{M-b+1}$ of length b computed from the posterior samples $\{p_k^{(i)}\}_{k=1}^M$. We compare it to an estimate of σ_i^2 using sample variance $\hat{\sigma}_i^2$ of the raw chain, $\{p_k^{(i)}\}_{k=1}^M$. Formally, we approximate the effective sample size of the length M MCMC chain as

$$M_{eff}^{(i)} = M \frac{\hat{\sigma}_i^2}{\hat{C}_i}.$$

Table 3: The effective sample sizes for a chain of length $M = 1000$ regarding all 16 probabilities for possible vector outcomes of visibles. BwTNLV would require at least 4.7 times as many MCMC iterations to achieve the same amount of effective information about the posterior distribution.

Outcome	BwTNLV	BwTNML	Outcome	BwTNLV	BwTNML
1	73.00	509.43	9	83.47	394.90
2	65.05	472.51	10	95.39	327.35
3	87.10	1229.39	11	70.74	356.56
4	72.64	577.73	12	81.40	338.30
5	71.67	452.01	13	105.98	373.59
6	66.49	389.78	14	132.61	306.91
7	84.30	660.37	15	82.15	365.30
8	75.46	515.09	16	98.05	304.57

The effective sample sizes for a chain of length $M = 1000$ for inference about each of the $2^4 = 16$ model probabilities are presented in Table 3. These range from 304.57 to 1229.39

for BwTNML, while BwTNLV only yields between 65.05 and 132.61 effective draws. Thus, BwTNLV would require at least 4.7 times as many iterations to be run of the MCMC chain in order to achieve the same amount of effective information about the posterior distribution. For this reason, consistent with the ACF results in Figure 12, BwTNLV does not seem to be an effective method for fitting the RBM if computing resources are at all limited.

Figure 13 shows the posterior probability of each possible $\mathbf{v} \in \{-1, 1\}^4$ after fitting the RBM model in the two ways detailed in this section (excluding BwTNLV). The black vertical lines show the true probabilities of each image based on the parameters used to generate the training set while the red vertical lines show the empirical distribution for the training set of 5,000 vectors. From these posterior predictive checks, it is evident that BwTNML produces the best fit to the data. However, this method requires a marginalization step to obtain the probability function of visible observations alone, which is infeasible for a model with n_H of any real size.

5 Discussion

RBM models constitute an interesting class of undirected graphical models that are thought to be useful for supervised learning tasks. However, when viewed as generative statistical models, RBMs are prone to forms of model impropriety such as near-degeneracy, S-instability, and uninterpretability. Additionally, these models are difficult to fit using a rigorous methodology, due to the dimension of the parameter space coupled with the size of the latent variable space.

In this paper, we have presented three honest Bayes MCMC-based methods for fitting RBMs. Common practice is to use a kind of MCMC to overcome fitting complexities. However because of the size of the space to be filled with MCMC iterates, convergence and mixing of these methods will be slow. Marginalization over the latent variables in the model can improve mixing, but is numerically intractable due to the necessity of repeated calculation of the normalizing constant.

Ultimately, it is not clear that RBM models are useful as generative models. Due to the

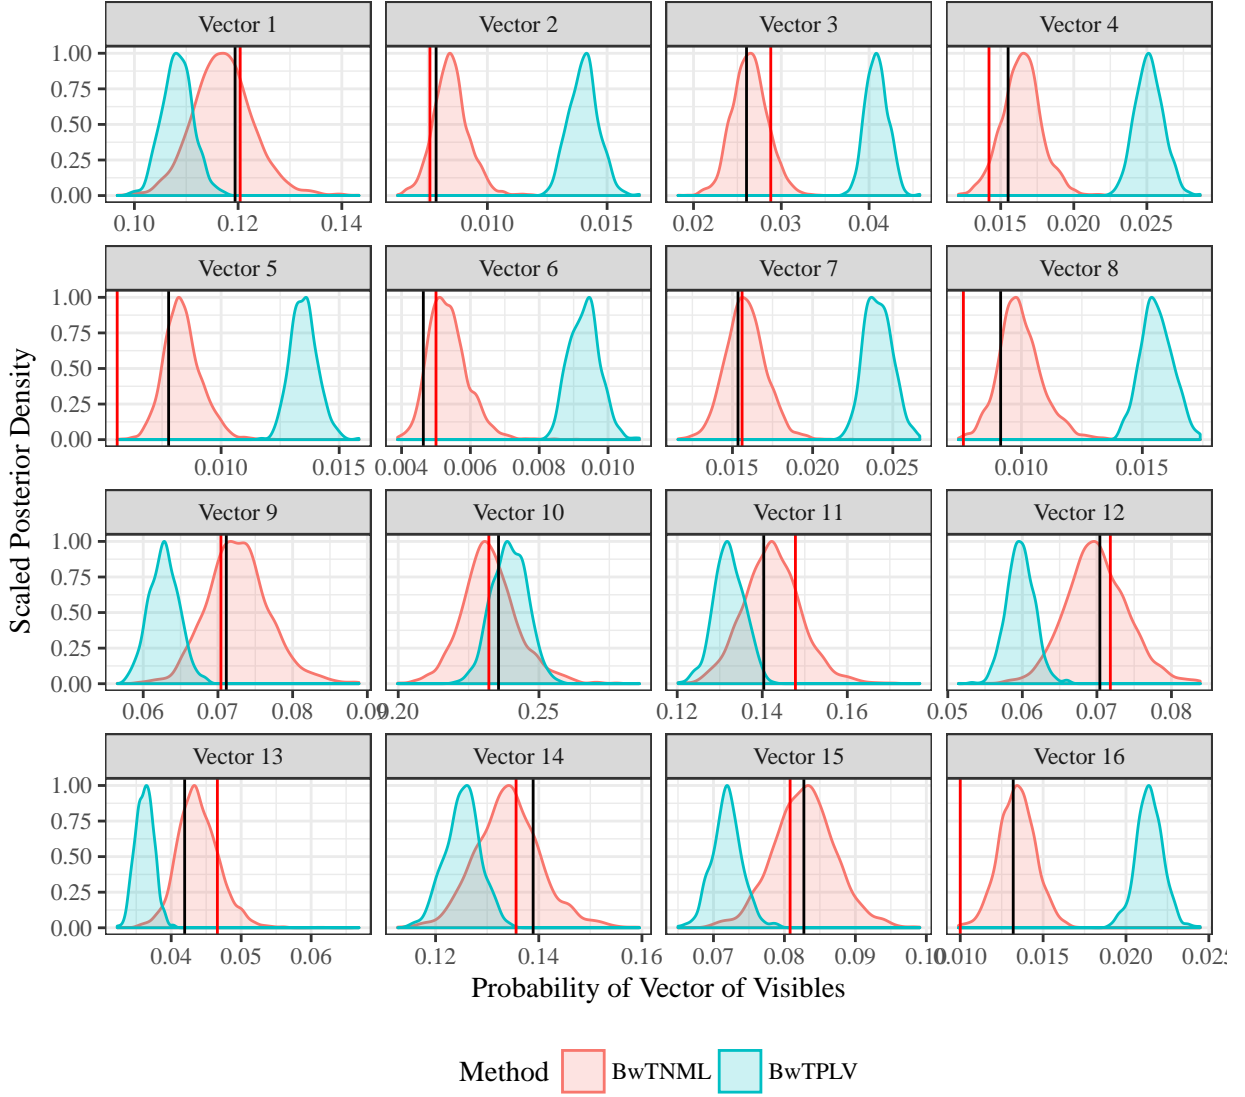


Figure 13: Posterior probabilities of $16 = 2^4$ possible realizations of 4 visibles using two of the three Bayesian fitting techniques, BwTPLV and BwTNML. The black vertical lines show the true probabilities of each vector of visibles based on the parameters used to generate the training data while the red vertical lines show the empirical distribution. BwTNML produces the best fit for the data, however is also the most computationally intensive and least feasible with a real dataset.

extreme flexibility in this model class, rigorous likelihood-based fitting for a RBM will typically merely return the (discrete) empirical distribution for visibles. Even if a rigorous likelihood-based fitting method is practically possible, we know what it will produce for fitted probabilities before we even begin. For these reasons, we are skeptical of the claims made about RBMs as generative tools.

References

- Fisher, R. A. 1922. “On the Mathematical Foundations of Theoretical Statistics.” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 222 (594-604). The Royal Society: 309–68.
- G. E. P. Box, W. J. Hill. 1967. “Discrimination Among Mechanistic Models.” *Technometrics* 9 (1). [Taylor & Francis, Ltd., American Statistical Association, American Society for Quality]: 57–71.
- Gelman, Andrew, Kenneth Shirley, and others. 2011. “Inference from Simulations and Monitoring Convergence.” *Handbook of Markov Chain Monte Carlo*, 163–74.
- Handcock, Mark S. 2003. “Assessing Degeneracy in Statistical Models of Social Networks.” Center for Statistics; the Social Sciences, University of Washington. <http://www.csss.washington.edu/>.
- Hinton, Geoffrey. 2010. “A Practical Guide to Training Restricted Boltzmann Machines.” *Momentum* 9 (1): 926.
- Hinton, Geoffrey E. 2002. “Training Products of Experts by Minimizing Contrastive Divergence.” *Neural Computation* 14 (8). MIT Press: 1771–1800.
- Hinton, Geoffrey E, Simon Osindero, and Yee-Whye Teh. 2006. “A Fast Learning Algorithm for Deep Belief Nets.” *Neural Computation* 18 (7). MIT Press: 1527–54.
- Kaiser, Mark S. 2007. “Statistical Dependence in Markov Random Field Models.” *Statistics Preprints* Paper 57. Digital Repository @ Iowa State University. <http://lib.dr.iastate.edu/>

[stat_las_preprints/57/](http://stat.las.preprints/57/).

Kaplan, Andee, Daniel Nordman, and Stephen Vardeman. 2016. “A Note on the Instability and Degeneracy of Deep Learning Models.” *In Preparation*.

Larochelle, Hugo, and Yoshua Bengio. 2008. “Classification Using Discriminative Restricted Boltzmann Machines.” In *Proceedings of the 25th International Conference on Machine Learning*, 536–43. ACM.

Le Roux, Nicolas, and Yoshua Bengio. 2008. “Representational Power of Restricted Boltzmann Machines and Deep Belief Networks.” *Neural Computation* 20 (6). MIT Press: 1631–49.

Lehmann, E. L. 1990. “Model Specification: The Views of Fisher and Neyman, and Later Developments.” *Statistical Science* 5 (2). Institute of Mathematical Statistics: 160–68.

Li, Jing. 2014. “Biclustering Methods and a Bayesian Approach to Fitting Boltzmann Machines in Statistical Learning.” PhD thesis, Iowa State University; Graduate Theses; Dissertations. <http://lib.dr.iastate.edu/etd/14173/>.

Montufar, Guido, and Nihat Ay. 2011. “Refinements of Universal Approximation Results for Deep Belief Networks and Restricted Boltzmann Machines.” *Neural Computation* 23 (5). MIT Press: 1306–19.

Montúfar, Guido F, Johannes Rauh, and Nihat Ay. 2011. “Expressive Power and Approximation Errors of Restricted Boltzmann Machines.” In *Advances in Neural Information Processing Systems*, 415–23. NIPS.

Nguyen, Anh Mai, Jason Yosinski, and Jeff Clune. 2014. “Deep Neural Networks Are Easily Fooled: High Confidence Predictions for Unrecognizable Images.” *ArXiv Preprint ArXiv:1412.1897*. <http://arxiv.org/abs/1412.1897>.

Salakhutdinov, Ruslan, and Geoffrey Hinton. 2012. “An Efficient Learning Procedure for Deep Boltzmann Machines.” *Neural Computation* 24 (8). MIT Press: 1967–2006.

Salakhutdinov, Ruslan, and Geoffrey E Hinton. 2009. “Deep Boltzmann Machines.” In *International Conference on Artificial Intelligence and Statistics*, 448–55. AI & Statistics.

Schweinberger, Michael. 2011. “Instability, Sensitivity, and Degeneracy of Discrete Ex-

ponential Families.” *Journal of the American Statistical Association* 106 (496). Taylor & Francis: 1361–70.

Smolensky, Paul. 1986. “Information Processing in Dynamical Systems: Foundations of Harmony Theory.” DTIC Document.

Srivastava, Nitish, and Ruslan R Salakhutdinov. 2012. “Multimodal Learning with Deep Boltzmann Machines.” In *Advances in Neural Information Processing Systems*, 2222–30. NIPS.

Srivastava, Nitish, Ruslan R Salakhutdinov, and Geoffrey E Hinton. 2013. “Modeling Documents with Deep Boltzmann Machines.” *ArXiv Preprint ArXiv:1309.6865*. <http://arxiv.org/abs/1309.6865>.

Szegedy, Christian, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2013. “Intriguing Properties of Neural Networks.” *ArXiv Preprint ArXiv:1312.6199*. <http://arxiv.org/abs/1312.6199>.

Zhou, Wen. 2014. “Some Bayesian and Multivariate Analysis Methods in Statistical Machine Learning and Applications.” PhD thesis, Iowa State University; Graduate Theses; Dissertations. <http://lib.dr.iastate.edu/etd/13816/>.